

A Reply to Hahn and Wallsten[†]

James Bessen*
Research on Innovation and Boston University (Visiting Researcher)

Robert M. Hunt**
Federal Reserve Bank of Philadelphia

March 10, 2004

* jbessen@researchoninnovation.org

** Ten Independence Mall, Philadelphia, PA 19106. Phone: (215) 574-3806. Email:
bob.hunt@phil.frb.org

[†]: The views expressed here are those of the authors and do not necessarily represent the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System.

A Reply to Hahn and Wallsten (2003)

In the recent European debate over software patents, fourteen economists issued a letter to the European Parliament arguing against expanding patent protection to software inventions without restriction (Anderson et al., 2003). They released this letter with a report (David et al., 2003) which cited our working paper, “An Empirical Look at Software Patents.” In that paper, we did not find that granting software patents has increased R&D. In fact, we found evidence that firms may be substituting software patents for R&D. These results seem inconsistent with the traditional incentive theory of patents.

In December of 2003, Robert Hahn and Scott Wallsten issued a critique of our paper (Hahn and Wallsten, 2003). This note is a response to their criticisms. But first, we wish to emphasize a point of agreement between these two papers (and the fourteen economists): Much more research should be done on the effects of granting software patents. Hahn and Wallsten argue that a policy question should not be decided on the basis of one article, and we heartily agree. This leads to an important question for policy makers in Europe: Do they currently have sufficient data to make an informed judgment about extending patent protection to computer programs?

We begin by presenting a synopsis of the main criticisms in Hahn and Wallsten and our response, which is based in part on our latest revision of the paper:

1. *Hahn and Wallsten argue that our definition of “software patent” is not rigorous.*

We compare our definition of “software patent” to others in the literature. We find that our definition is reasonably accurate and is more accurate (in terms of false positives and false negatives) than a definition based on the International Patent Classification (IPC) system. The empirical results are robust to the substitution of an IPC based definition for our own.

2. *Hahn and Wallsten argue that our sample is not representative.*

Our sample spans those firms responsible for the vast majority of patents and nearly all R&D reported by U.S. firms contained in Standard & Poor’s *Compustat* data set. While our coverage of the newest and smallest firms is not as good, our tests for sample selection bias suggest there is very little or no effect on results.

3. *Hahn & Wallsten criticize our analysis of whether software patents and R&D are complements or substitutes.*

Their criticism seems to be based on a misunderstanding of the concept of substitute and complementary inputs. Our analysis uses a simple modification of a textbook econometric test for complementary inputs and our results are robust to different specifications.

4. *Hahn & Wallsten argue that we ignore the problem of endogeneity in our regressions*

We certainly do not ignore this potential. We test for the possibility of endogeneity bias and find none. We also report estimates from an instrumental variables regression.

5. *Hahn and Wallsten argue that “even when firms engage in strategic patenting, it is not possible to determine whether that is good or bad in a welfare sense,” in particular because the availability of patents may reduce barriers to entry.*

We test for an asymmetric response of young or small firms to the availability of software patents and generally find no difference. If there is a positive social return to software patenting despite the apparent negative effects on R&D investments, it has yet to be articulated in either theoretical or empirical research.

Detailed Discussion

Criticisms of Our Data

- Hahn and Wallsten say: *B&H[Bessen and Hunt] use an extremely broad definition of “software patent,” provide no analytical justification of their definition, and ignore the IPC classification system, which other authors have used to construct more rigorous definitions of “software patents.”*

Economists have used two main approaches to classifying patents into industries or technologies (Griliches, 1990): to read and individually classify the patents, or to use functional classifications assigned by the patent office.

Our approach is to classify a sample of patents by reading them and, by trial-and-error, to develop a search algorithm that attempts to replicate this classification. Then, to validate the accuracy of the algorithm, we test it against a random sample of patents that we have also read and classified. Using a random sample of 400 patents, this algorithm had a false positive rate of 16% (that is, 16% of the patents the algorithm said were software patents, were not) and a false negative rate of 22% (that is, it failed to identify 22% of the patents we categorized as software patents). We also tested it against random samples classified by independent researchers (although they used a more restrictive definition, see paper) and found a false negative rate of 8% and a false positive rate less than an upper bound of 26%.

Hahn and Wallsten describe our classification as “extremely broad” based on the fact that they discovered two errors. Our classification does, of course, make errors. However, it is unlikely that our patent counts are very far off because our false positive and false negative rates are roughly equal. More important, tests show (below) that there is no *systematic* bias in our patent counts that might influence our results.

Hahn and Wallsten claim that a classification based on the IPC classes, such as one used by Graham and Mowery (2003), is more rigorous. We believe that while an IPC based definition may be sufficient for the kind of analysis that Graham and Mowery (2003) perform, our definition is superior and better suited for the purposes of our analysis.

Patent classes have well-known problems for classifying patents into technologies. Griliches (1990) argued that the patent classification system is “based primarily on technological and functional principles and is only rarely related to economists’ notions of products or well-defined industries (which may be a mirage anyway). A subclass dealing with the dispensing of liquids contains both a patent for a water pistol and for a holy water dispenser. Another subclass relating to the dispensing of solids contains patents on both manure spreaders and toothpaste tubes.” Consider an example more closely related to the subject at hand—a class for “visual representation of data” includes software, computer displays, and electrical and mechanical signs that pre-date electronic computers.

We compared the IPC based definition in Graham and Mowery (2003) to our own. We found that, based on our random sample of 400 patents, this definition had a higher false positive rate and a much higher false negative rate (30% and 74% respectively). This definition would exclude half of the patents obtained by the largest publicly-traded software firms and a majority of the pure software and Internet patents identified in Allison and Tiller (2003). In contrast, our definition accounts for about 4/5 of all patents obtained by the top 200 software firms in the 1990s. Nevertheless, we also ran our main tabular and regression analyses using a definition based on the IPC classifications identified in Graham and Mowery (2003) instead of our own definition. *All* of the main empirical results still hold, although standard errors are uniformly higher. This suggests that there is no systematic bias introduced by our classification method relative to Graham and Mowery’s.

- Hahn and Wallsten say: *The authors match firms to their R&D and financial data only [sic] if those firms existed and were publicly-traded in 1989, assuming they used the NBER match file. As a result, their analysis excludes all firms that went public in the 1990s.*

In the revised paper, the matched sample has been substantially expanded. It contains 1,230 subsidiaries and 305 firms not contained in the NBER match file. The matched sample covers 91% of the deflated R&D and 89% of the deflated sales of R&D-performing firms in Compustat over the years 1980-99. It also covers 68% of the patents granted to domestic non-governmental organizations (mostly firms) and 73% of software patents granted to these organizations. These coverage ratios are roughly constant over the entire twenty year period.

Our matched sample is thus highly representative of the firms that perform most of the R&D and obtain most of the patents. It is less representative of small and entrant firms, however, although it does contain substantial numbers of these firms. To make sure that sample selection does not bias our regression results, the paper reports a series of Heckman sample selection model estimates. In these we find either that we can reject the hypothesis of sample selection bias or any such bias has relatively little effect on the estimated coefficients.

- Hahn and Wallsten say: *B&H use patent grants rather than applications, contrary to nearly all literature. Grants follow applications by about two years, on average. This lag means that the analysis captures almost no information regarding patenting behavior after 1996, which is the period of greatest interest.*

In the revised paper, we date patents by application date and uses data through 1997. This covers the critical years following the *Alapat* decision (1994). We were not surprised to find qualitatively the same results as reported for patents dated by grant year.

- Hahn and Wallsten say: *The authors find that manufacturing firms receive most of the software patents, as well as most patents overall, but do not mention that they actually hold a smaller share of software patents than their share of total patents would suggest.*

This is a misleading argument. There is no particular reason to suggest that industries that acquire a large portion of patents on other technologies should also acquire a large portion of software patents, especially if they do not do much software development. The real issue is to analyze the sources of the growth in software patenting (next).

Are Software Patents “Cheaper” Than Other Patents?

In the revised paper, we replace the discussion of relative value with a closely-related discussion of patent propensity.¹ Building on the model used by Hall and Ziedonis (2001), we estimate patent propensity. That is, we measure the extent to which firms obtain software patents after controlling for differences in investment in R&D and physical capital, other firm characteristics, and the employment of programmers and engineers across industries.

We find that industries and firms known for strategic patenting (computers, electronics, instruments and IBM) have a *much greater* propensity to obtain software patents than firms in other industries. They also have a larger propensity to patent in general. Moreover, as Hall and Ziedonis (2001) and Hall (2003) show, these industries have increased their patent propensities dramatically since the creation of the Court of Appeals for the Federal Circuit in the mid-80s. This is consistent with the view that strategic patenting is driving this increase.

We also find a rapid increase in patent propensity over time that cannot be explained by increases in programmer productivity or any of the other factors for which we can control.² We conclude that the cost effectiveness of patents has increased dramatically—that is, patents became less costly to obtain and/or they delivered greater appropriability.

Hall and Ziedonis found that entrant firms had a higher patent propensity and they interpreted this as partial evidence that patents may help entrant firms obtain financing (see also Hall, 2003). We do not find a similar effect in our regressions for software patents. In fact, entrant software firms are substantially *less likely* to obtain software patents.

Finally, Hahn and Wallsten suggest that because software patents have more claims and receive more subsequent citations, that these are “valuable” patents. This is consistent with other research (Allison et al., 2003) that finds that software patents are more likely to be litigated. But this is evidence of *private* value, not *social* value. In particular, if software patents are used to build strategic patent portfolios, they may be privately valuable because they serve to block competitors, yet this may have no redeeming social value.

¹ We plan to explore the question of relative value in a separate paper so we do not address those points here.

² In the latest version of the paper, we use a variety of techniques to control for unobserved firm heterogeneity. The qualitative results do not change.

Are R&D and Software Patents Complements or Substitutes?

This is a key investigation of our paper. We find that in the 1990s software patents appear to substitute for R&D and that this effect is concentrated in the strategic patenting industries identified in the previous literature. This raises some doubt about the relevance of the traditional “incentive theory” for software patents.

- Hahn and Wallsten say: *One of the authors’ results in the “cost” analysis reveals a statistically significant positive correlation between patents and R&D, suggesting that they may be complements.*

This is an erroneous inference. Suppose that coffee and tea are substitutes; that is as the *price* of one rises, consumption of the other rises. Now some households will consume more of both (e.g., wealthy households) and some will consume less of both (e.g., poorer households). In the data, it is entirely possible that the quantity of coffee and tea consumed is positively correlated across households, but we would not conclude from this data that they are complementary goods.

Now a decrease in the price of tea relative to the price of coffee will lead households to consume less coffee *relative* to the amount of tea they consume. Wealthy households may still consume more coffee than poor households, but less than they consumed earlier. The substitution effect can be observed more clearly by examining share of expenditures spent on coffee—if coffee’s share of total expenditures decreases with a drop in the price of tea, then coffee and tea cannot be complements.

In our paper, we argue that the cost effectiveness of software patents—the appropriability conferred per dollar cost of obtaining them—increased over time. This almost certainly is true in absolute terms, but for our purposes all that matters is the change in the cost effectiveness of software patents relative to all other patents. This is in effect a reduction in the relative price of software patents. We don’t observe this price directly, but we do observe firm behavior that is correlated with it—changes over time in the ratio of software patents obtained by a firm to the total number of patents obtained by that firm. Hence we use changes in this ratio—*software patent share*—as an (inverse) measure of the change in relative prices.

Now large companies tend to have both more R&D and more patents than smaller ones. This tells us nothing about the relationship between these two inputs in the production of rents earned by the firm. But if R&D and software patents are complementary inputs in the production of those rents, we should expect that firms who focus increasingly on software patents will also devote a larger share of their expenditures to R&D. All else equal, under the traditional incentive

theory of patents, we should expect these two variables to be positively correlated, or at worst uncorrelated. But, as we show in the paper, during the 1990s changes in software patent share are negatively correlated with changes in R&D as a share of expenditures (costs) or revenues (sales). That suggests that firms substituted software patents for R&D in the 1990s.

- Hahn and Wallsten say: *Patents, R&D, and sales affect each other, but the authors do not account for these endogenous relationships, making it difficult to interpret coefficients.*

Hahn and Wallsten argue that because the number of software patents obtained is an endogenous variable, we should estimate a simultaneous equations model as in Arora et al. (2003). In principal, one would like to specify an equation that explains R&D and patenting decisions as a function of each other, as well as other variables, and find some way of plausibly identifying the resulting supply and demand effects. The problem is in the identification strategy—what factors affect patenting, but not R&D?

We opt for a very different approach from that used in Arora et al. (2003). We modify a standard textbook technique for estimating the elasticity of substitution for factors of production. This approach uses input cost shares as the left hand side variable and input prices as right hand side variables (see for example Greene, 1997, Chapter 15, and Berndt, 1990 chapter 9).³ A typical equation would have the ratio of R&D spending to costs on the left hand side, and a variety of input prices on the right hand side. When industry-level prices are used, they are reasonably exogenous to the firm. In other words, a firm would not expect its decisions about R&D spending to affect the input prices it faces. In that case estimating multiple equations is not required.⁴

Our regressions are run in differences (see Table 6 of the revised paper) and we use changes in the software share of patents as a measure of changes in the cost effectiveness of software patents relative to other patents. Hahn and Wallsten are correct that this variable may be correlated with the error terms, as it is at least potentially endogenous. To test for this, we performed the same regressions using instrumental variables—we instrumented the share of software patents with one and/or two year lagged values (in differences). This did not change our results and a Hausman test could not reject the hypothesis that our coefficients were unbiased. In other words, we adequately control for the simultaneous determination of software patents and R&D and it did not affect our results.

³ The cost of using this approach is that we are no longer estimating causal relationships, but rather correlations.

⁴ Typically, multiple cost share equations are estimated together as part of a Seemingly Unrelated Regression. This increases efficiency, however, each individual equation estimated alone is nevertheless consistent.

- Hahn and Wallsten say: *B&H use the ratio of R&D to sales as the dependent variable rather than the level of R&D, and the share of patents that are software rather than the number of software patents as the relevant independent variable. As a result, it is not possible to determine whether any coefficient estimate implies that R&D and patents are complements or substitutes.*

In a textbook economy, one could infer something about the relationship between two goods by observing changes in the quantities consumed as prices are changed. But in a real economy, demand and supply curves shift. For this reason, the most common technique for estimating elasticities of substitution uses *shares* rather than *levels*. Hahn and Wallsten propose another estimation, but they don't explain how they would control for changing supply and demand or why they feel that the commonly-used method is inappropriate.

As we explain in the paper, we favor estimating a regression based on cost or revenue shares, as it implicitly controls for changes in demand and other effects. In addition, the resulting coefficients that can be translated into cross price elasticities of factor demand. The only wrinkle in an otherwise orthodox estimation strategy is our use of software patent share as a proxy for the unobserved relative price of software patents. Other specifications are possible. If instead we run our basic regression using the change in the log of software patents and the change in the log of other patents as explanatory variables, the coefficient of interest is again negative and highly significant.⁵ The paper reviews the many other ways we tested the robustness of this result.

- Hahn and Wallsten say: *B&H eliminate firms where R&D is greater than one-half sales. While these firms may be outliers, the paper contains no discussion of this choice, which effectively truncates the dependent variable and thus potentially biases the results.*

In the revised paper, we initially estimate the R&D regression with *no* trimming. Then, to overcome some data limitations, we use a slightly different dependent variable (R&D to sales), but, to avoid spurious effects, we exclude observations where R&D exceeds one-half sales. This makes little difference in the key coefficient, only a small number of observations are excluded by the trimming (mainly biotech startups), and we report results from a Heckman sample selection model that finds no evidence of bias.

Conclusion

A basic finding of our paper is that those firms that acquired more software patents during the 90s (relative to their total patenting) reduced their R&D intensity relative to other firms. This result is

⁵ The coefficient on the change in the log of other patents is not statistically significant.

difficult to reconcile with the claim that software patents increased incentives for performing R&D. We suggest that results would be consistent with strategic patenting behavior, a theory also supported by some of our other findings.

Hahn and Wallsten argue we are pursuing the wrong line of research and raise a laundry list of errors in our working paper. We believe that we have demonstrated the robustness of our empirical results and the soundness of our model, which is, after all, a simple modification of a standard technique for determining whether factors are complements or substitutes.

Hahn and Wallsten also argue that even if software patents encouraged strategic patenting, they may still have had beneficial effects for small and startup firms. This ignores our evidence that new software firms acquire relatively few software patents, and that the substitution effect occurs among small firms as well as large firms. In this regard, software patents may have a different effect than patents acquired by semiconductor design firms.

They also argue that “even when firms engage in strategic patenting, it is not possible to determine whether that is good or bad in a welfare sense.” They suggest that strategic patenting could facilitate technology transfer or knowledge spillovers. Although *cross-licensing* may facilitate knowledge spillovers (if the cross-licenses include real technology transfer rather than merely the more common agreement not to litigate), we are unaware of any research suggesting that *strategic portfolio behavior* encourages knowledge spillovers. Moreover, if firms obtain software patents for reasons other than to protect their innovations—i.e. they are strategic patentees—welfare could be reduced (see Bessen 2003).

Hahn and Wallsten issued their critique of our paper as a response to the economists’ letter to the EU parliament. Hahn and Wallsten admit they have taken the unusual step of criticising a *working* paper because they are afraid it will “unduly influence policy decisions.” They conclude with a call for more research. This is of course the same conclusion *already reached* in the economists’ letter. But those economists go further, suggesting it might be prudent for the European Parliament not to act until more is understood about the economic effects of software patents. Given the paucity of empirical evidence, this does not seem to be unreasonable advice...

References

Allison, John R. and Mark A. Lemley. 2000. “Who’s Patenting What? An Empirical Exploration of Patent Prosecution,” *Vanderbilt Law Review*, Vol. 58, pp. 2099-2148.

- Allison, John R. and Emerson H. Tiller, 2003, "Internet Business Method Patents," in Wesley M. Cohen and Stephen A. Merrill, eds., *Patents in the Knowledge-Based Economy*, National Research Council, Washington: National Academies Press, p. 259.
- Ashish Arora, Marco Ceccagnoli, and Wesley M. Cohen. 2003. "R&D and the Patent Premium," NBER Working Paper No. 9431.
- Bessen, James. 2003. "Patent Thickets: Strategic Patenting of Complex Technologies," ROI Working Paper.
- Bessen, James and Robert Hunt. 2004. "An Empirical Look at Software Patents," Federal Reserve Bank of Philadelphia Working Paper 03-17R, March 2004.
- Graham, Stuart J. H. and David C. Mowery, 2003. "Intellectual Property Protection in the U. S. Software Industry," in Wesley M. Cohen and Stephen A. Merrill, eds., *Patents in the Knowledge-Based Economy*, National Research Council, Washington: National Academies Press, pp. 219-58.
- Greene, William H. 1997. *Econometric Analysis*, 3rd edition. Upper Saddle River, N.J.: Prentice Hall.
- Griliches, Zvi. 1990. "Patent Statistics as Economic Indicators: A Survey," *Journal of Economic Literature*, Vol. 28, pp. 1661-1707.
- Hahn, Robert W. and Scott Wallsten. 2003. "A Review of Bessen and Hunt's Analysis of Software Patents," AEI-Brookings Joint Center, mimeo.
- Hall, Bronwyn H. 2003. "Exploring the Patent Explosion," invited lecture at the ZEW Workshop on Empirical Economics of Innovation and Patenting, Mannheim, Germany, March 14/15, 2003.
- Hall, Bronwyn H., Adam B. Jaffe, and Manuel Trajtenberg. 2001a. "The NBER Patent Citations Data File: Lessons, Insights and Methodological Tools," NBER Working Paper No. 8498.
- Hall, Bronwyn H., and Rosemary Ham Ziedonis. 2001. "The Patent Paradox Revisited: An Empirical Study of Patenting in the U.S. Semiconductor Industry, 1979-1995" *RAND Journal of Economics*. Vol. 32, pp. 101-128.